

CSL overview for ACEs

Ben Kilminster, Jedong Lee, Kevin McFarland, Tony Vaiciulis
University of Rochester

Hiroyuki Matsunaga, Makoto Shimojima, Hirokazu Kobayashi
University of Tsukuba

April 10, 2003

- CSL description
- monitoring

CSL ACE Help Page

Much CSL information is documented on CSL ACE help page

Data Taking



[Ops Plan](#) , [Run Plan](#) , [Good Run Definition](#)



[\[Run Control\]](#) – [\[Error Handler\]](#)



[Important CDF DAQ Processes \(PROCMON\)](#)



[TevMon Instructions](#)



[SVX for ACEs](#)



[Software Event Builder ACE Help Page](#)



[Online Computers](#)



[CDF Electronic Logbooks \(e-logs\)](#)



[Data Acquisition \(DAQ\) system](#)



[L3/EVB Manual for ACEs](#)



[L3 Manager Information](#)



[CSL ACE Help Page](#)
(data to consumers/
disks / look area)



[Starting Consumers](#)
[Running event display offline](#)



[VxWorks nodes basics](#)
(includes explanation
of front-end errors)

CSL ACE Help Page

[About the
CSL
Starting/Stoppin
Look area
Monitoring
Calibration
CSL
Troubleshooting
Call CSL
Expert
CSL Home
Page](#)

Consumer-Server/Logger ACE Help Page

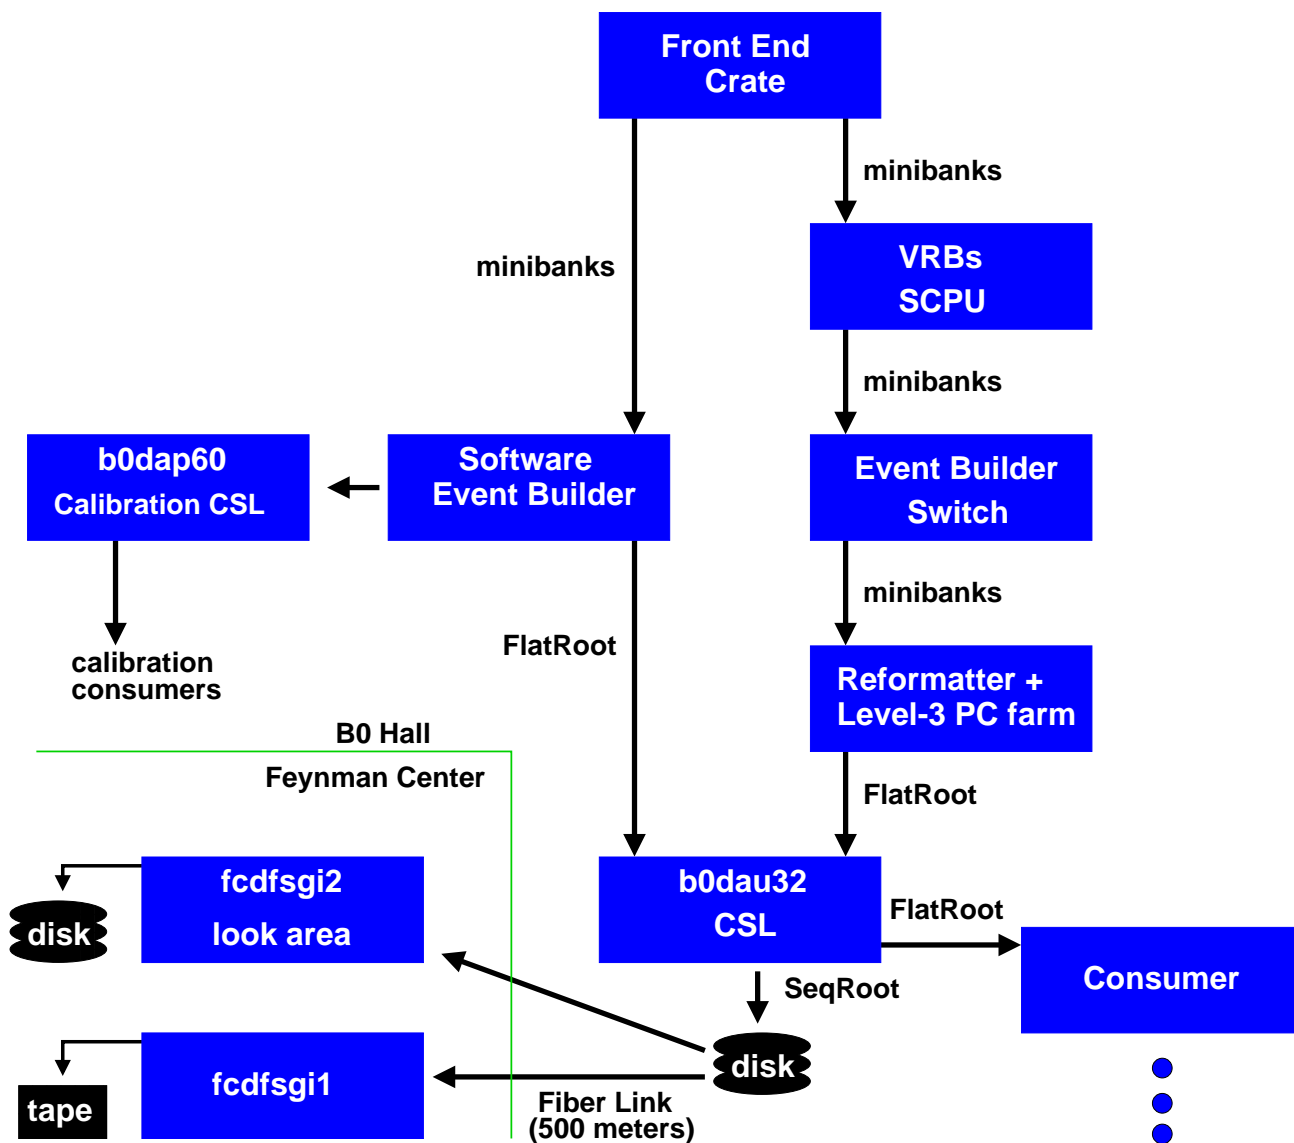
About the CSL

Questions, comments, suggestions? Send mail to
cdf_csl@fnal.gov

The Consumer-Server/Logger (CSL) is the central online hub through which all CDF data taken during Run II must pass. Events passing the Level-3 trigger are written to disk and served to consumers by the CSL. The raw data files on disk are then moved to the Feynman Computing Center (FCC) for further processing. Consumers are online monitoring processes, spying on a fraction of the data passing through the CSL. The diagram below shows how the CSL fits into the overall dataflow.



CSL description



primary CSL functions

- receive events from Level-3 PC farm at 20 MB/sec ($75 \text{ Hz} \times 250 \text{ kB/event}$)
- receive events from the software event builder
- write events to disk at 20 MB/sec
- handle as many consumer requests as possible (5-10 MB/s total)

Starting/stopping the CSL

Page a CSL expert before starting or stopping the CSL.

During normal running, the CSL never needs to be started or stopped. It is always "on", ready to receive events.

Should it be advised by an expert :

The shift crew is able to start and stop the CSL.

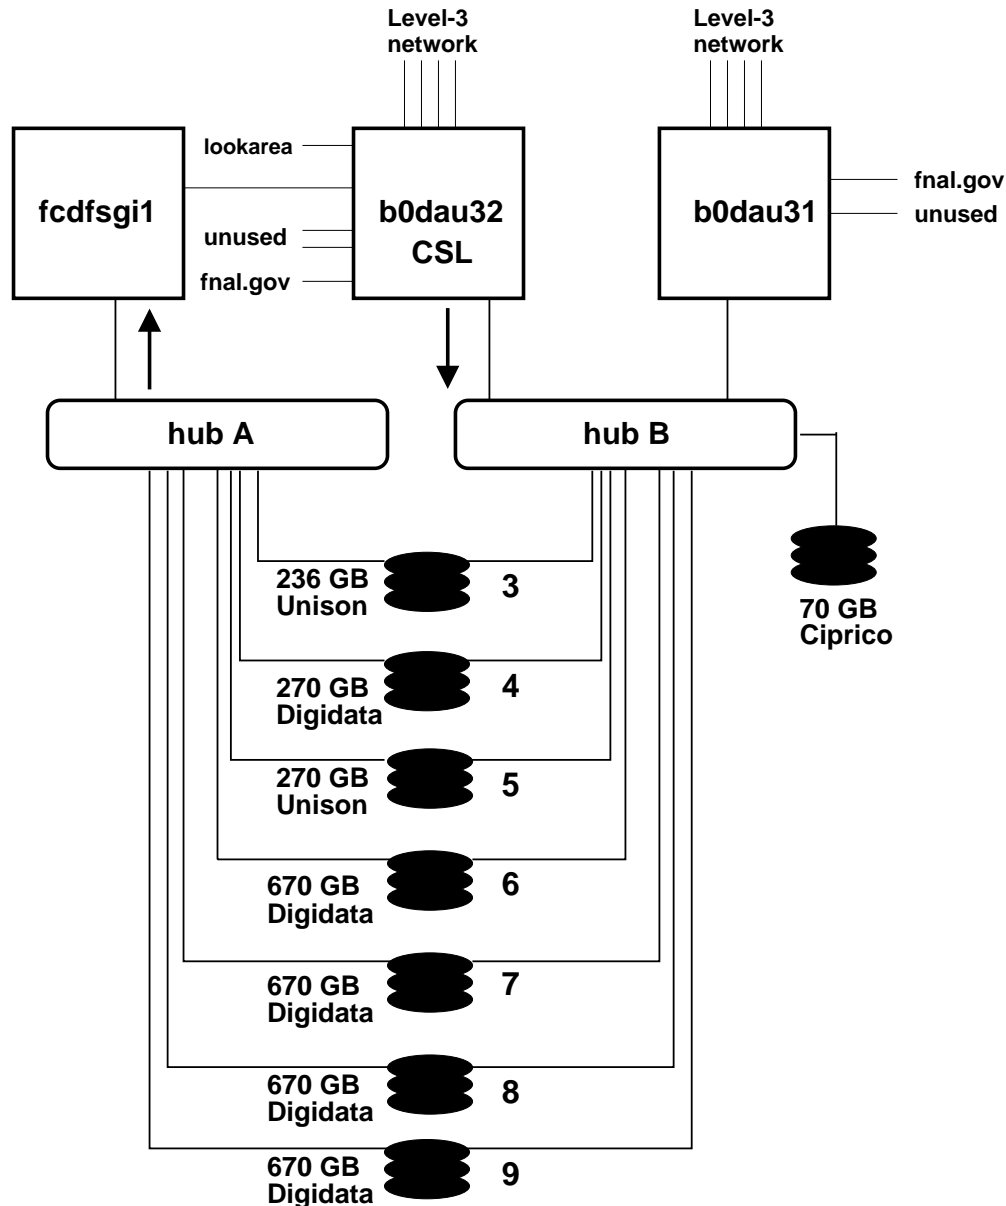
Log on to b0dau32 as user cdfdaq and type the following commands:

b0dau32: cslcom check (check if CSL processes are running)

b0dau32: cslcom stop (stop all CSL processes)

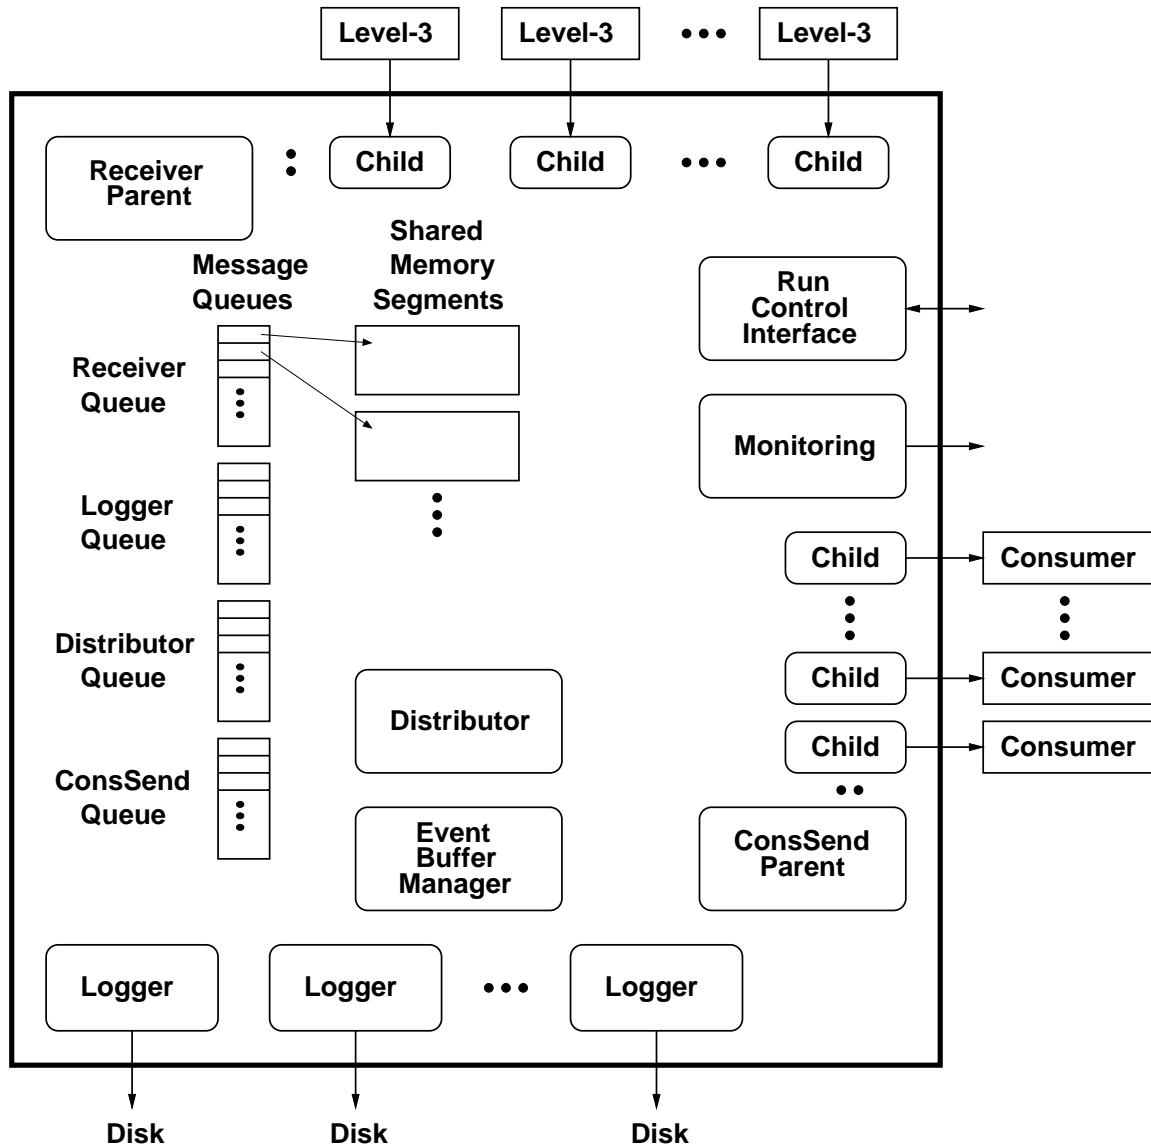
b0dau32: cslcom start (start all CSL processes)

CSL hardware overview



- b0dau32, an SGI 2200 Server
(4 CPUs, large I/O bandwidth)
located on third floor of B0
- ~3.2 TB of disk space on third floor of B0 (7 RAIDs)
- dual ported disks allow both fcdfsgi1 in FCC and b0dau32 to access disks

CSL software: overview



The CSL is a “server”. Possible “clients” include

- Level-3 output node processes
- software event builder processes
- 24 hour sender in partition 14
- consumers

Events are stored in shared memory buffers. Flow of events between processes inside CSL achieved by means of message queues.

CSL software: monitoring

The monitoring process collects CSL status information and sends it via a smartsockets message to the CSL monitoring display.

There are two kinds of CSL monitoring:

- the CSL display (snapshot of current CSL state)
- the CSL history plots (show rates, disk space, events logged, etc. versus time)

Using CSL display to check current CSL status:

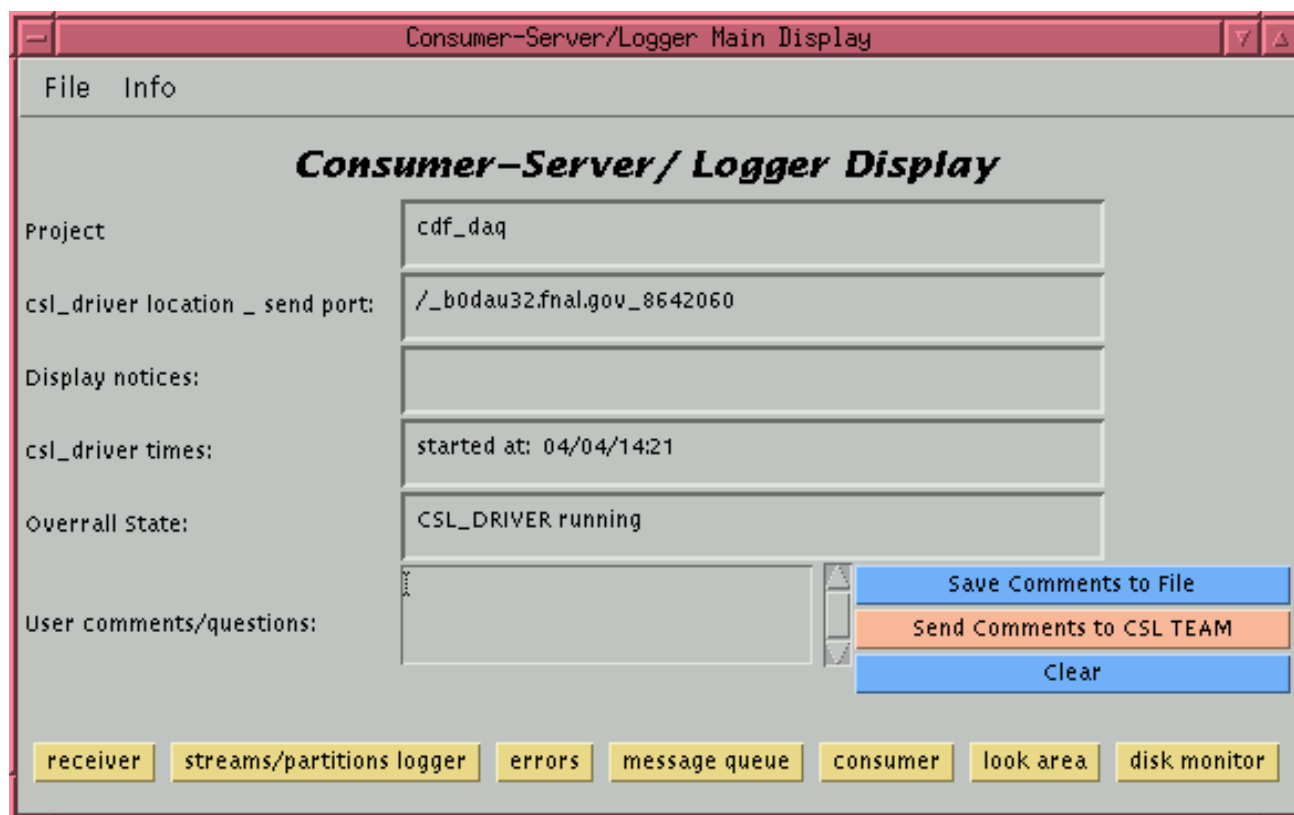
login to machine on online cluster

setup fer devel

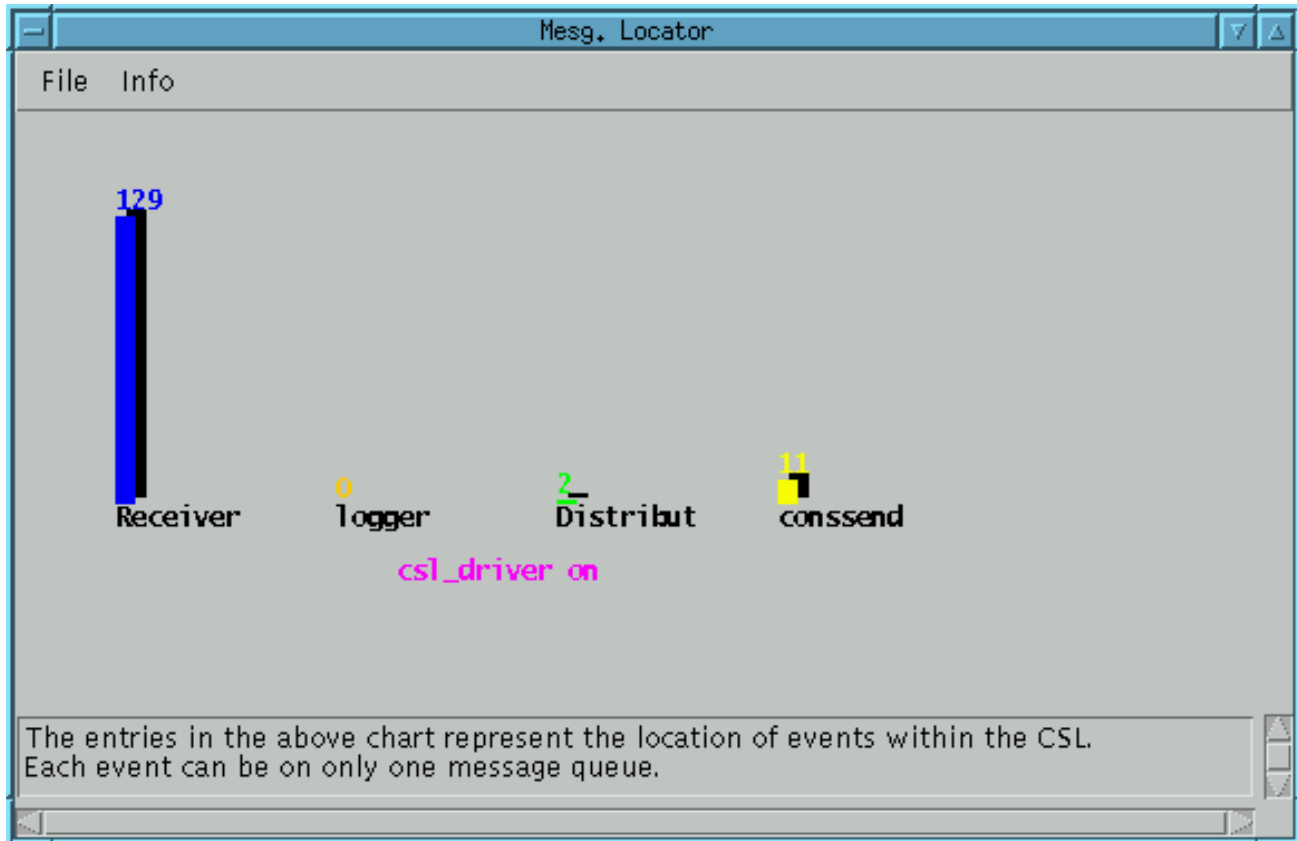
daqmon (or cslmon for just the csl display)

A GUI should appear on your screen. Click on CSL button.

Use project name cdf_daq (this should be the default). The CSL monitor main display window should appear.



CSL monitoring: message queues



There are 150 internal buffers in the CSL. Each buffer can store one event. The display shows four message queues. Each message on a queue points to a buffer where an event may be stored.

Receiver (left, blue) queue shows how many buffers are free to store new events from Level-3. If this is ALWAYS zero and "logger" queue has all the buffers then there is a problem. Contact CSL expert.

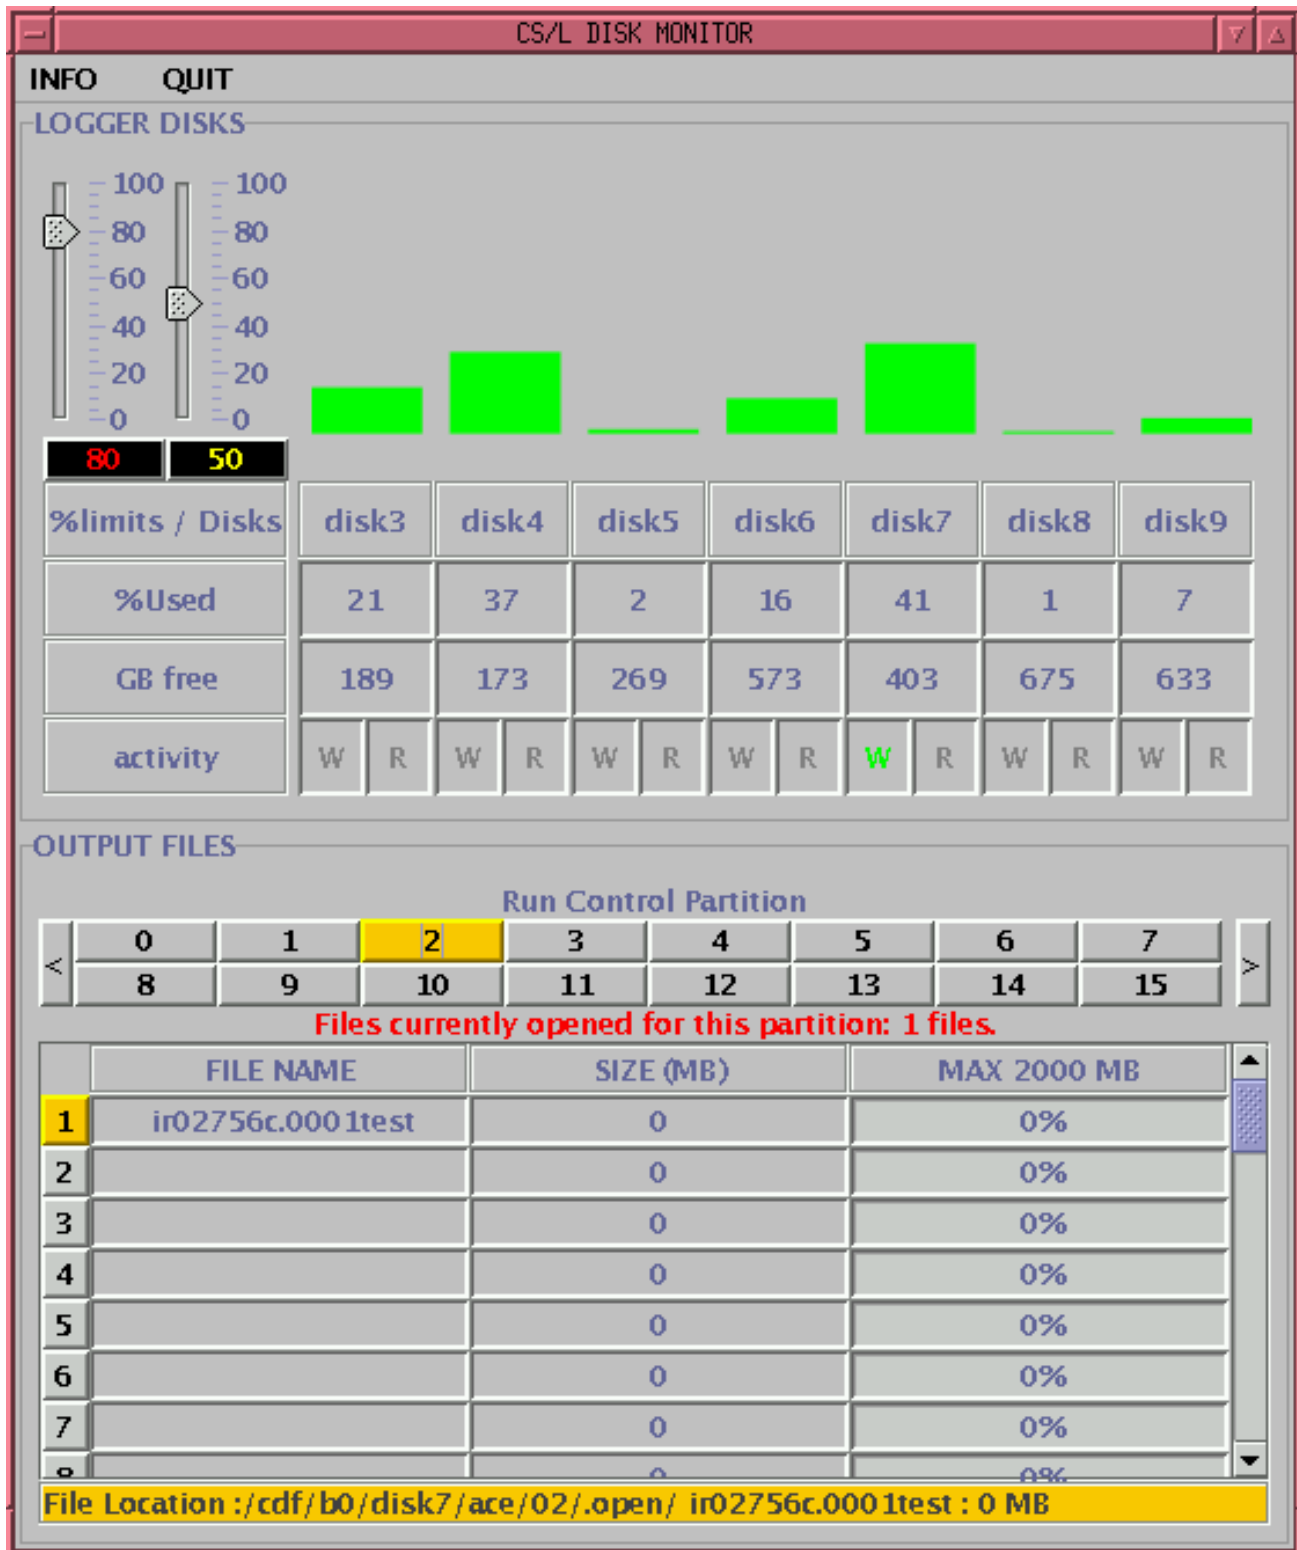
CSL monitoring: receiving

[illegible]

For each client sending events to the CSL, there is a receiver process.

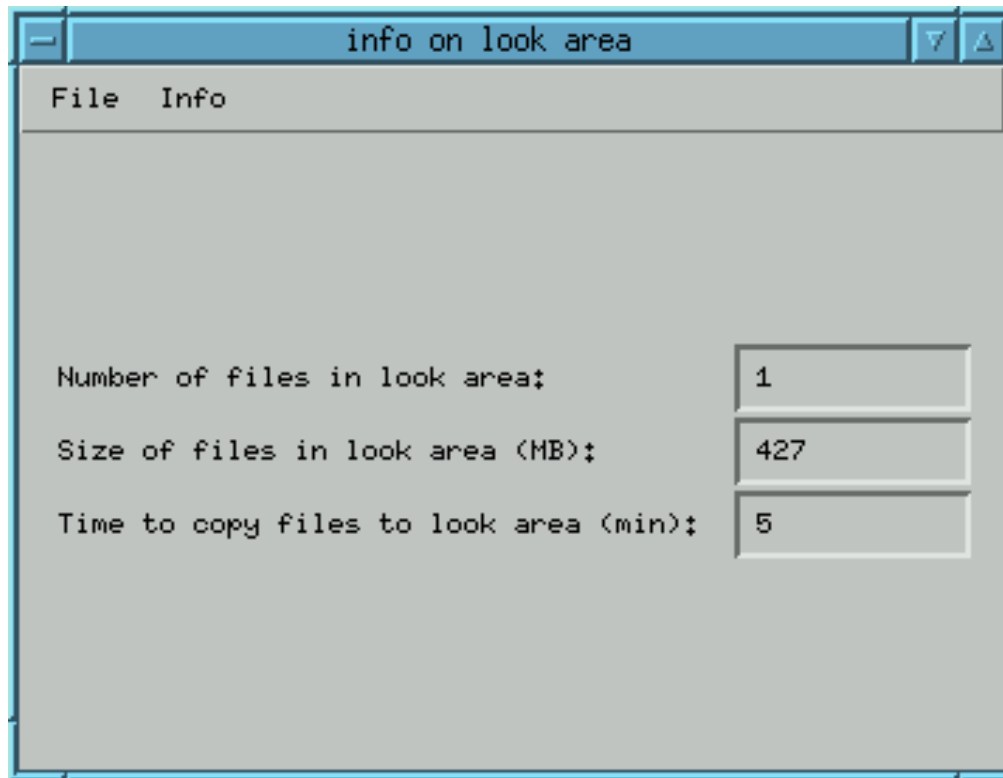
The CSL monitoring GUI displays information for each receiver process: client node name, partition, number of events received, and average and instant rates (event rate, event size, activity rate).

CSL monitoring: logging



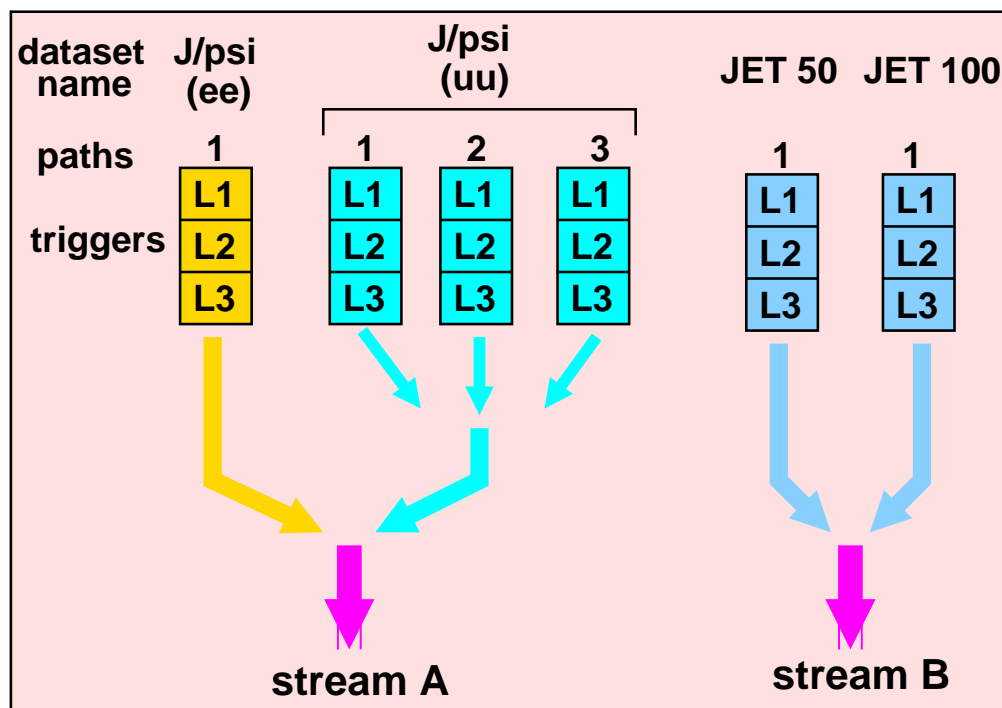
How full is each CSL disk? Are output files growing larger?

Look Area



- Default is to copy all Stream A physics files, the first few files from some other physics, test, cosmic, and calibration streams.
- Ace Controls
 - Aces can change from default to ALL or NONE under advice of expert
 - However ... be careful about low bandwidth for file copying and swamping look area with junk

Paths, datasets, streams in Run 2



path: AND of Level-1, Level-2, Level-3 triggers.

dataset: OR of all paths defined for that dataset.

stream: collection of datasets

Events coming out of Level-3 are “streamed”: tagged as belonging to a particular stream.

- CDF DAQ system can run in multi-partition mode
- each partition is independent of other partitions
- CSL writes events into separate files for different partitions
- CSL writes events corresponding to different streams within a partition into separate files

Run II: 10 streams, 50-100 datasets

CSL monitoring: serving consumers

File
Info

Consumer		Statistics			Status					
Cons. Type	Cons. id / part. id	Host name or address	Evts Req.	Events Recvd	Inst Rate	Conn	Waiting	Active	Comple	
TrigMc	7902 / C	kCcap58 fnal cov	1	0	0					
(MB/s) for Current Consumers (C.C.)	C	Consumer children ending in error				46				
Current Events Requested for C.C.	1									
Combined Events Requested for C.C.	1									
Events Sent to C.C.	C	<div> <div>Show average rate</div> <div>Show most recent rate</div> </div>								
Events Sent to ALL Consumers										

CSL LAW #1: write ALL events to disk

CSL LAW #2: send events to consumers as long as it does not break LAW #1


Consumers do not see all events (maybe 5-10 MB/sec). Consumers can request events by partition number, stream, L1/L2/L3 triggers.

CSL history plots

Consumer–Server/Logger Monitoring Plots

CSL specific plots		
receiver, logger, consumer rates	24 hours	one week
number of events logged	24 hours and one week	
number of processes	24 hours	one week
CPU usage: receivers,loggers,consends	24 hours	one week
CPU usage: driver,distributor,bufman,monsend	24 hours	one week
general b0dau32 plots		
disk space usage	24 hours	one week
load average	24 hours and one week	
free memory	24 hours and one week	
CPU usage: global	24 hours	one week

CSL history plots: disk usage



Red = total disk space used

Gray/blue = subset of disk space used for files which are waiting to be written to tape

The difference between red and gray = files which are not going to be written to tape. They are automatically deleted if more disk space is needed.

An empty plot indicates the disk is not visible for some reason.

disk4 = stream A

disk5 = stream B

disk6 = unused

disk7 = stream J,I


disk8 = stream C,D,G

disk9 = stream E,H

disk3 = overflow from other disks if needed

control disk = log files

CSL history plots: logging rate



Calibration CSL

A special version of the CSL software runs on b0dap60. Useful for some calibration runs which require

- guaranteed delivery of all events to the consumer OR
- a very large event size
(expected Run 2 event size is about 250 kB, main CSL can accept up to 3 MB, calibration CSL can accept up to 17 MB)

Calibration CSL does not log any data to disk.

Data File Catalog

CSL writes information into the Data File Catalog database for each output data file:

- file size
- total number of events in the file
- run number
- first event number in the file
- last event number in the file
- run section numbers

This must be done before the files are put onto tape.

Troubleshooting

Here are some reasons why you might think there is a problem with the CSL

- The CSL does not acknowledge a run control transition.
- There is an error message from the CSL in the Error Logger.
- You believe Level-3 or the software event builder is sending events to the CSL and one or more of the following is true
 - the consumers are not receiving any events
 - the events do not seem to be written to disk by the CSL
 - no files for the runs you are taking appear in the "look" area on fcdfsgi2

Some things you should check before paging a CSL expert:

- Is the CSL receiving any events?
- If the CSL is not receiving events it may or may not be a CSL problem.
- Is the CSL sending events to consumers?
- Is the CSL writing events to disk?
- Did the CSL send an error message to the Error Logger?

CSL ACE web pages describes how to answer these questions.

AFTER going through the checklist, if you still think there is a problem with the CSL, then page CSL expert. Four person rotation now: Ben Kilminster, Jedong Lee, Hiro Matsunaga, Tony Vaiciulis